

AMENDMENTS TO THE SPECIFICATION

Please replace the paragraph beginning at page 3, line 1 with the following rewritten paragraph:

-- What is needed is a system and method for controlling distribution of e-mail messages (or other network communications, collectively "e-mail" messages) that ~~reduces~~ reduce the burden on network resources of recipients and intermediaries and/or allows for distribution of e-mail messages in a prioritized manner according to preferences. Prioritization and preferences allow a recipient or intermediary to define criteria and/or set preferences to regain control of the utilization of the recipient's or intermediary's system resources.--

Please replace the paragraph beginning at page 3, line 9 with the following rewritten paragraph:

--The present invention allows recipients, ISPs, ESPs, and other network communication (message) recipients to control how their network's or ~~systems~~ system's resources, such as network connectivity bandwidth, are used and/or allocated for use to distribute messages from others. The present invention also allows for preferential prioritized treatment of compliant messages sent from trusted senders, and lower priority treatment of non-compliant messages from non-compliant and/or irresponsible senders. The present invention further provides for lessening or avoiding the impact of virus, spam, and denial of service attacks, and the ability to load balance incoming messages onto a cluster of servers.--

Please replace the paragraph beginning at page 3, line 18 with the following rewritten paragraph:

--The present invention provides flow prioritization and spam squelcher functionality whereby network communications (messages) either carry priority information, or are assigned priority information based on a shared characteristic with other messages. The priority information is used to determine how, when and/or whether to deliver or process the message, e.g. by delaying the message for a fixed time, routing to a "junk" folder, or deleting. Preferences for receipt of messages by priority level may be ~~communication~~ communicated to upstream hosts along a network path. Accordingly, an intermediary host may reject, delay and/or delete messages that the intended recipient does not wish to receive or the intermediary does not wish to process (e.g. messages carrying viruses, fraudulent messages, or spam messages). This pushes the burden of low-priority messages back to the sender, thereby reducing or eliminating burdens on network/system resources of the recipient and/or intermediaries between the recipient and the sender. Accordingly, it can "squelch" spam messages at or close to their source. Trusted senders complying with prescribed practices may include priority information allowing for delivery of their messages with higher priority.--

Please replace the paragraph beginning at page 8, line 6 with the following rewritten paragraph:

--A priority level for the incoming e-mail message is then identified, as shown at step 14 of Figure 1. This may be performed in a variety of ways. In one embodiment, the e-mail messages are configured to carry priority information in accordance with the present invention. For example, such priority information may include a priority flag, such as "Immediate Action Required," "For Reference Only," "For Information Only," "Highest Priority," "Highest Priority," etc., or may include a

categorization of the e-mail message, e.g. "Internal Business Communication," "External Business Communication," "Personal," "Friends and Family," "Bills and Statements," etc. It should be noted that higher priority treatment may be sold to or purchased to from senders as desired to obtain higher priority treatment of their messages. In one embodiment, the message may include a message-type specifier such as a keyword or a seal graphic/image object (indicating compliance with certain laws, regulations and/or best practices standards) contained in the viewable portion of the message, e.g., body, subject line, etc., of the message, that may be used as the priority information. Alternatively, the priority information may be a compiled header that is a datastring, such as an alphanumeric or ASCII character string, stored in a special field of the message's header information that includes information indicating a priority level, or information that may be used to determine a priority level. The compiled header may include encrypted or formatted data that may be understood by the network appliance, and preferably not by others. Such header information is typically hidden from the recipient. Provision of such header information, message specifiers and/or other priority information is disclosed in commonly assigned U.S. Application No. 09/793,263, filed February 26, 2001, (Attorney Docket No. P24528 USA), U.S. Application No. 09/793,296, filed February 26, 2001 (Attorney Docket No. P24,618 USA) and U.S. Application No. 09/792,936, filed February 26, 2001 (Attorney Docket No. P24,773 USA), the disclosures of which are hereby incorporated herein by reference. The compiled header is generated by appropriate software and/or an appropriate appliance at the sender or an intermediary, such as an ISP or ESP. The header includes information indicating that it is a message that is compliant with the present invention and/or a certain standard associated with the present invention, e.g. identifying the sender of the

message as a "trusted sender." This header is generated programmatically from the full header information contained generally in all e-mail messages (generated at the sender), e.g. at the same time by the same software or a component thereof.--

A4
Please replace the paragraph beginning at page 11, line 7 with the following rewritten paragraph:

AS
-- For incoming e-mail messages that do not carry priority information, e.g. do not include a specifier, compiled header information, etc., other information must be used to identify a priority level for the message in step 14. In such a case, other information typically carried by an e-mail message, or associated with an e-mail message, is used in an inventive way in accordance with the present invention to identify a priority level. For example, information that may be used includes a network address of a sending or intermediary (e.g. ISP) system, a sub-network of a sending or intermediary system, a domain name of a sending or intermediary system, or other network path information typically found in network communication headers, such as TCP/IP packet headers, or application path information such as network path information of systems previously processing and/or routing the message, sender identity domain or other information typically carried by an e-mail header (e.g., SMTP header information such as TO, FROM, SUBJECT, DATE, mail agent, "received" history, etc.), or other geographic origin (as determined by known techniques), source, origin, path or other information that may be shared by multiple messages and thereby provide identifying information by which future messages from the same geographic origin, source, origin, path, etc. may be prioritized, etc.--

Please replace the paragraph beginning at page 12, line 6 with the following rewritten paragraph:

-- Figure 2 is a flow diagram 20 of an exemplary method for using a heuristic approach for identifying a priority level for an e-mail message's network path according to the present invention. Accordingly, the method of Figure 2 may be used as part of step 14 of Figure 1 when an e-mail message does not carry priority information. Conceptually, the method of Figure 2 identifies characteristic information that may be shared by multiple messages and performs a statistical analysis on those messages to determine which messages are likely to be from irresponsible senders or be undesirable, etc., and then assigns all messages having that shared characteristic information a corresponding priority level. Accordingly, messages having a shared characteristic are assigned a priority level based on some metric or other shared information. For example, messages delivered from a certain sender along a certain network path may be sampled to determine levels of messages containing viruses, to determine whether the messages are undeliverable (which is often the case when a dictionary attack spam method is used), or whether the messages are spam, e.g. as determined by pattern matching or other known techniques for detecting abusive messages.--

Please replace the paragraph beginning at page 13, line 1 with the following rewritten paragraph:

--Referring now to Figure 2, the exemplary method starts with identification of an incoming e-mail message's network path, as shown at steps 21 and 22. For example, such information may be ascertained by referring to TCP/IP packet headers. This is performed by the network appliance 100 of Figure 3, specifically, by

the heuristic engine 102. One or more pluralities of messages sharing a network path are then sampled, as shown at step 24. For example, the sampling rate of messages may be predetermined and static, e.g. 5% of all messages received on the network path. For example, this may be established by the system, system administrator, or recipient. The sampling rate may be determined ~~be~~ by reference to the rules engine 104 or another repository for storing system or recipient preferences.--

Please replace the paragraph beginning at page 14, line 4 with the following rewritten paragraph:

-- For the sample messages, e.g. 5% of messages arriving along a certain network path, a value for a sender metric is determined, as shown at step 26. For example, sender metrics may include a delivery success rate metric indicating the percentage of messages that are delivered (or undeliverable). For example, this may be achieved by maintaining a hash table of valid addresses in order to determine deliverability of sampled messages, or by proxy processing of e-mail "bounce" messages by the network appliance. It should be noted that this table of processing may be performed dynamically, in real time. Methods and techniques for doing so ~~is~~ are straightforward as will be appreciated by those skilled in the art. The rate of undeliverable messages is typically unusually high when the messages are sent using a dictionary attack or brute force spam method because many of the recipient addresses are merely guesses at valid network addresses. Alternatively, a spam rate metric may be determined to indicate the proportion of messages from a certain path that are deemed to be spam, and therefore undesirable. For example, messages may be determined to be spam using content-based analysis, such as

pattern matching, as is well known in the art. While this technique may not be completely effective for filtering, it is useful for heuristic analysis. As another alternative, a virus rate metric may be used to reflect the number of messages along a given network path that carry viruses, e.g. using virus-checking software. Any suitable metric and/or value may be used. The sampled messages are scanned by the e-mail scanners 106 of the network appliance 100 (Figure 3), and suitable software for implementing the e-mail scanners is well-known in the art, or requires straightforward modification of software well-known in the art.--

Please replace the paragraph beginning at page 15, line 21 with the following rewritten paragraph:

--Referring again to Figure 1, the rule base 104 is next referenced to determine a prescribed delay for the identified priority level. More generally, the rule base may specify any kind of delivery instructions for the identified priority level. For example, the rule base may store rules indicating that messages with a "HIGHEST" priority level are to be delivered directly to the recipient's mail server 140 without any unnecessary delay. For example, this may be performed without regard to network resource availability or load. Alternatively, a rule may specify that messages with a "LOWEST" priority be delayed, e.g. by holding the message for a fixed period of time, or until network resource availability reaches a certain level, etc. Similar rules may apply to any message characteristic, e.g. "BUSINESS." By delaying delivery, more network/system resources are available for distribution of higher priority messages, and distribution of the lower priority message will not likely use network resources that are need needed to deliver higher priority messages. By not delivering the message at the desired time, or by rejecting the message altogether

and not delivering the message, the burden may be shifted from the recipient's network resources to the senders (or an intermediary's) network resources, which then must resend the message. The technique of delaying the message in fact delays the entire connection, so no recipient mail server resources are used until that connection is processed. The benefit is realized over time as the total number/volume of connections from a certain path is contained to desired levels through the cumulative delay in processing of individual connections. An application of this concept is the balancing of paths in terms of a percentage of total available resources. This delaying or rejecting of network connections provides a mechanism for controlling the total number of connections for a particular message source or path.--

Please replace the paragraph beginning at page 17, line 13 with the following rewritten paragraph:

-- The delaying and/or eventual delivery of messages is performed by the e-mail connection processor 108 of the network appliance, as shown in Figure 3. The connection processor 108 receives incoming connections for delivery of e-mail messages and calls other components of the network appliance 100 as necessary. The connection processor 108 also creates network connections as necessary to deliver e-mail messages at the appropriate time, e.g. after the prescribed delay, to deliver messages in accordance with the prioritization techniques described herein. The connection processor may be implemented with software using network architecture layer 4 and/or layer 7 switching and/or load balancing techniques that are well known in the art. Modification of such techniques to provide for delayed delivery, etc. in accordance with the present invention is straightforward,

A10
as will be appreciated by those skilled in the art.--

Please replace the paragraph beginning at page 20, line 16 with the following rewritten paragraph:

A11
--If such a request is ignored by a poorly-behaved email host (as evidenced by no change in e-mail traffic volume), the inventive system initiates a Spam Squelch by limiting the number and volume of connections, minimally at a network level, from the offending host. For example, the network appliance tracks volume/number of inbound connections per host, and ~~slow slows~~ and/or ~~stop stops~~ the TCP and/or SMTP connection build process for squelched hosts. Additionally, when a source path is identified, e.g. for spam, it may be traced to an ISP of the spam sender. Accordingly, the ISP may be identified and pressured to demand compliance of the spam sender with a certain industry standard of best practices. The ISP may be held responsible for failure to obtain compliance of the spam sender by imposing a lower priority level to all messages receive via the ISP's communication path, thereby imposing a heavy burden on the ISP.--

Please replace the paragraph beginning at page 22, line 3 with the following rewritten paragraph:

A12
--Accordingly, the present invention allows the ISPs/ESPs and other email hosts ~~are able to control the volume of messages to within their network's/systems'~~ processing capability, and still respect the priority of compliant e-mail messages. Furthermore, the content of email messages is not of concern to the ISPs/ESPs, so longstanding arguments by the ISP community regarding their role as telecommunications infrastructure (with respect to Communications Decency Act

(CDA), COPPA, etc etc.) will not be threatened by their inspection of e-mail message content. In this manner, ISPs maintain their lack of responsibility for content on the grounds that they do not inspect content in any way, but rather merely deliver communications traffic. Accordingly, the present invention provides for spam reduction without the need for such content inspection.--

X¹²
Please replace the paragraph beginning at page 27, line 1 with the following rewritten paragraph:

--A method for controlling distribution of network communications (messages). An incoming message either carries priority information, or is assigned priority information based on a shared characteristic with other messages. The priority information is used to determine how and/or when to deliver the message, e.g. by delaying the message for a fixed time. Preferences for receipt of messages by priority level may be communication communicated to upstream hosts along a network path. Accordingly, an intermediary host may reject and/or delay messages that the intended recipient does not wish to receive. This pushes the burden of low-priority messages back to the sender, thereby reducing or eliminating burdens on network/system resources of the recipient and/or intermediaries between the recipient and the sender. Accordingly, it can "squelch" spam messages at or close to their source. Trusted senders complying with prescribed practices may include priority information allowing for delivery of their messages with higher priority.--
